

CLAIMS

1. A method for data access, comprising:
 - defining an ontology for application to a set of diverse data sources comprising data having predefined semantics;
 - 5 associating with the ontology one or more logical rules applicable to the semantics of the data in the data sources;
 - receiving a query from a user regarding the data;
 - determining a query plan for responding to the query by selecting one or more of the data sources responsively to the ontology and by identifying an operation to be applied to the data responsively to the applicable logical rules; and
 - 10 generating a response to the query in accordance with the query plan.
2. The method according to claim 1, wherein the logical rules comprise a validation rule, and wherein the query plan comprises validating the data from at least one of the data sources responsively to the validation rule.
- 15 3. The method according to claim 1, wherein the logical rules comprise a mapping rule, such that at least one of the data sources is mapped to the ontology in accordance with the mapping rule, and wherein the query plan comprises determining an applicability of the at least one of the data sources to the query responsively to the mapping rule.
4. The method according to claim 1, wherein the logical rules comprise a joining rule, and
20 wherein the query plan comprises selecting a key responsively to the joining rule, and joining the data from two or more of the data sources using the key.
5. The method according to claim 4, wherein selecting the key comprises analyzing the data so as to select one or more fields in the two or more of the data sources for use as the key so as to provide a desired statistical probability that the data will be joined correctly.
- 25 6. The method according to claim 1, wherein the logical rules comprise a transformation rule, and wherein the query plan comprises transforming the data in at least one of the data sources from a first value that is held in the at least one of the data sources to a second value responsively to the transformation rule.

7. The method according to claim 1, wherein the logical rules comprise a business logic rule, and wherein the query plan comprises processing the data from at least one of the sources responsively to the business logic rule.

8. . . . The method according to claim 1, wherein the logical rules comprise an access rule, 5 and wherein the query plan comprises selecting at least one of the data sources for use in generating the response responsively to the access rule as applied to the user who submitted the query.

9. The method according to any of claims 1-8, wherein defining the ontology comprises 10 associating a respective wrapper with each of the data sources, so as to transform the data from each of the data sources from a native format to an ontological format determined by the ontology, and wherein generating the response comprises applying the operation using the wrapper, and then reporting the data from the wrapper to a hub that links the data sources following application of the operation.

10. The method according to claim 9, wherein the operation applied by the wrapper 15 comprises joining the data from two or more of the data sources.

11. The method according to claim 9, wherein the operation applied by the wrapper comprises mapping values of the data.

12. The method according to claim 11, wherein mapping the values comprises normalizing the data from a native representation to an ontological representation.

20 13. The method according to claim 9, wherein the query plan comprises a group of sub-queries, and wherein generating the response comprises sending the sub-queries from an agent running on the hub to respective wrappers of a plurality of the data sources, and combining the data reported from the wrappers in order to produce the response.

14. The method according to claim 13, wherein sending the sub-queries comprises 25 invoking two or more of the wrappers to operate in parallel.

15. The method according to claim 9, wherein associating the respective wrapper comprises distributing an advertisement of each of the data sources in accordance with the ontology, and wherein determining the query plan comprises discovering each of the data

sources responsively to the advertisement, and building the query plan based on the discovered data sources.

16. The method according to claim 9, wherein reporting the data comprises sending data packets over a network, the packets comprising semantic content in a form determined by the

5 ontology, and upon receipt of the data packets at the hub, verifying legitimacy of the packets responsively to the semantic content.

17. The method according to claim 9, wherein reporting the data comprises streaming the data from the wrapper to a specified storage location.

18. The method according to claim 9, wherein reporting the data comprises moving the
10 data in a block operation from the wrapper to a specified storage location.

19. The method according to any of claims 1-8, wherein determining the query plan comprises collecting information regarding a topology and performance characteristics of the data sources, and selecting, responsively to the information, the data sources to be used in responding to the query.

15 20. A method for providing a user with access to a set of diverse information resources, which are configured to provide information with predefined semantics, the method comprising:

defining an ontology for application to the information provided by the set of diverse information resources;

20 associating with the ontology one or more logical rules applicable to the semantics of the information;

receiving a request from the user to access the information;

determining a plan for responding to the request by selecting one or more of the information resources responsively to the ontology and by identifying an operation to be applied to the information responsively to the applicable logical rules; and
25

generating a response to the request in accordance with the plan.

21. The method according to claim 20, wherein the request comprises a query for data held by the information resources.

22. The method according to claim 20, wherein the request comprises a subscription request, and wherein generating the response comprises providing a succession of responses over a period of time responsively to the subscription request.
23. The method according to claim 20, wherein the information resources are configured to provide Web services, and wherein the request specifies one or more of the Web services to be provided to the user.
24. The method according to claim 20, wherein the request specifies data to be written to a data repository associated with one or more of the information resources.
25. The method according to claim 20, wherein the request specifies a transaction to be carried out and recorded by one or more of the information resources.
26. The method according to any of claims 20-25, wherein the request specifies a process to be carried out by one or more of the information resources.
27. The method according to claim 26, wherein the request specifies an event, and wherein generating the response comprises carrying out the specified process responsively to an occurrence of the event.
28. A method for data access, comprising:
defining an ontology for application to a set of diverse data sources comprising data;
defining data access rights with respect to the ontology; and
controlling user access to the data responsively to the ontology of the data and the access rights applicable thereto.
29. The method according to claim 28, wherein defining the ontology comprises specifying a user ontology, and wherein defining the data access rights comprises assigning a classification to a user according to the user ontology, and wherein controlling the user access comprises comparing the classification to the access rights applicable to the data.
30. The method according to claim 29, wherein the diverse data sources are distributed among a set of autonomous organizations comprising at least first and second organizations, and wherein assigning the classification comprises classifying the user according to an organizational affiliation of the user so as to control access by users in the first organization to the data sources held by the second organization.

31. The method according to any of claims 28-30, wherein controlling the user access comprises receiving a query from a user to access the data in the data sources, determining a query plan for responding to the query by selecting one or more of the data sources responsively to the ontology such that the access rights permit the user to access the data in the
5 one or more of the data sources, and generating a response to the query in accordance with the query plan.

32. A method for data access, comprising:
defining a set of data resources providing access to data;
collecting information regarding a topology and performance characteristics of the data
10 resources;
receiving a query from a user regarding the data;
determining a query plan responsively to the query and to the information regarding the topology and performance characteristics; and
generating a response to the query in accordance with the query plan.

15 33. The method according to claim 32, wherein collecting the information comprises measuring respective load levels of two or more of the data resources, and wherein determining the query plan comprises selecting one of the data resources so as to balance the load levels.

20 34. The method according to claim 32 or 33, wherein the data resources are distributed among a set of autonomous organizations comprising at least first and second organizations, wherein the user submitting the query belongs to the first organization, and wherein determining the query plan comprises selecting, responsively to the performance characteristics, one of the data resources of the second organization for use in responding to the query.

25 35. A method for exchange of information, comprising:
establishing a virtual private network among a plurality of nodes, comprising at least first and second nodes, which are configured to communicate with one another over an underlying public physical network;
defining a semantic communication model for conveying data packets among the nodes
30 in the virtual private network, responsively to an ontology of the information;

sending a data packet over the virtual private network from the first node to the second node; and

filtering the data packet against the semantic communication model at the second node, so as to verify that the data packet is legitimate.

5 36. The method according to claim 35, wherein defining the semantic communication model comprises defining a closed set of semantic messages that may be carried by data packets in the virtual private network.

37. The method according to claim 35, wherein the nodes are distributed among a set of autonomous organizations.

10 38. The method according to any of claims 35-37, wherein the nodes comprise gateway nodes, which are configured to communicate with clients and data sources using native data formats, and to translate the native data formats to the semantic communication model for communication over the virtual private network.

15 39. The method according to claim 38, wherein the nodes further comprises hub nodes, and wherein establishing the virtual private network comprises configuring the gateway nodes to communicate with the hub nodes using the semantic communication model.

20 40. Apparatus for data access, comprising a hub processor, which is adapted to receive a definition of an ontology for application to a set of diverse data sources comprising data having predefined semantics, and to associate with the ontology one or more logical rules applicable to the semantics of the data in the data sources, and which is further adapted, upon receiving a query from a user regarding the data, to determine a query plan for responding to the query by selecting one or more of the data sources responsively to the ontology and by identifying an operation to be applied to the data responsively to the applicable logical rules, and to generate a response to the query in accordance with the query plan.

25 41. The apparatus according to claim 40, wherein the logical rules comprise a validation rule, and wherein the query plan comprises validating the data from at least one of the data sources responsively to the validation rule.

42. The apparatus according to claim 40, wherein the logical rules comprise a mapping rule, such that at least one of the data sources is mapped to the ontology in accordance with the

mapping rule, and wherein the query plan comprises determining an applicability of the at least one of the data sources to the query responsively to the mapping rule.

43. The apparatus according to claim 40, wherein the logical rules comprise a joining rule, and wherein the query plan comprises selecting a key responsively to the joining rule, and
5 joining the data from two or more of the data sources using the key.

44. The apparatus according to claim 43, wherein selecting the key comprises analyzing the data so as to select one or more fields in the two or more of the data sources for use as the key so as to provide a desired statistical probability that the data will be joined correctly.

45. The apparatus according to claim 40, wherein the logical rules comprise a
10 transformation rule, and wherein the query plan comprises transforming the data in at least one of the data sources from a first value that is held in the at least one of the data sources to a second value responsively to the transformation rule.

46. The apparatus according to claim 40, wherein the logical rules comprise a business logic rule, and wherein the query plan comprises processing the data from at least one of the
15 sources responsively to the business logic rule.

47. The apparatus according to claim 40, wherein the logical rules comprise an access rule, and wherein the query plan comprises selecting at least one of the data sources for use in generating the response responsively to the access rule as applied to the user who submitted the query.

20 48. The apparatus according to any of claims 40-47, and comprising at least one gateway processor, which is coupled to communicate with the hub processor, and which is adapted to associate a respective wrapper with each of the data sources, so as to transform the data from each of the data sources from a native format to an ontological format determined by the ontology, wherein the at least one gateway processor is adapted to apply the operation to the
25 data using the wrapper, and to report the data to the hub processor following application of the operation.

49. The apparatus according to claim 48, wherein the operation applied by the wrapper comprises joining the data from two or more of the data sources.

50. The apparatus according to claim 48, wherein the operation applied by the wrapper comprises mapping values of the data.

51. The apparatus according to claim 50, wherein mapping the values comprises normalizing the data from a native representation to an ontological representation.

5 52. The apparatus according to claim 48, wherein the query plan comprises a group of sub-queries, and wherein the hub processor is adapted to send the sub-queries to respective wrappers of a plurality of the data sources, and to combine the data reported from the wrappers in order to produce the response.

10 53. The apparatus according to claim 52, wherein the hub processor is adapted to invoke two or more of the wrappers to operate in parallel.

15 54. The apparatus according to claim 48, wherein the at least one gateway processor is adapted to distribute an advertisement of each of the data sources in accordance with the ontology, and wherein the hub processor is adapted to discover each of the data sources responsively to the advertisement, and to build the query plan based on the discovered data sources.

55. The apparatus according to claim 48, wherein the at least one gateway processor is adapted to report the data by sending data packets over a network, the packets comprising semantic content in a form determined by the ontology, and wherein the hub processor is adapted to verify legitimacy of the packets responsively to the semantic content.

20 56. The apparatus according to claim 48, wherein the at least one gateway processor is adapted to stream the data from the wrapper to a specified storage location.

57. The apparatus according to claim 48, wherein the at least one gateway processor is adapted to move the data in a block operation from the wrapper to a specified storage location.

25 58. The apparatus according to any of claims 40-47, wherein the hub processor is adapted to collect information regarding a topology and performance characteristics of the data sources, and to select, responsively to the information, the data sources to be used in responding to the query.

59. Apparatus for providing a user with access to a set of diverse information resources, which are configured to provide information with predefined semantics, the apparatus

comprising a hub processor, which is adapted to receive a definition of an ontology for application to the information provided by the set of diverse information resources and to associate with the ontology one or more logical rules applicable to the semantics of the information, and which is further adapted, upon receiving a request from the user to access the information, to determine a plan for responding to the request by selecting one or more of the information resources responsively to the ontology and by identifying an operation to be applied to the information responsively to the applicable logical rules, and to generate a response to the request in accordance with the plan.

5 60. The apparatus according to claim 59, wherein the request comprises a query for data held by the information resources.

10 61. The apparatus according to claim 59, wherein the request comprises a subscription request, and wherein the hub processor is adapted to provide a succession of responses over a period of time responsively to the subscription request.

15 62. The apparatus according to claim 59, wherein the information resources are configured to provide Web services, and wherein the request specifies one or more of the Web services to be provided to the user.

63. The apparatus according to claim 59, wherein the request specifies data to be written to a data repository associated with one or more of the information resources.

20 64. The apparatus according to claim 59, wherein the request specifies a transaction to be carried out and recorded by one or more of the information resources.

65. The apparatus according to any of claims 59-64, wherein the request specifies a process to be carried out by one or more of the information resources.

25 66. The apparatus according to claim 65, wherein the request specifies an event, and wherein the hub is adapted to cause the specified process to be carried out responsively to an occurrence of the event.

67. Apparatus for data access, comprising a hub processor, which is adapted to receive a definition of an ontology for application to a set of diverse data sources comprising data and a definition of data access rights with respect to the ontology, and which is adapted to control

user access to the data responsively to the ontology of the data and the access rights applicable thereto.

68. The apparatus according to claim 67, wherein the ontology comprises a user ontology, and wherein the hub processor is adapted to define the data access rights by assigning a classification to a user according to the user ontology, and to compare the classification to the access rights applicable to the data in order to control the user access.

69. The apparatus according to claim 68, wherein the diverse data sources are distributed among a set of autonomous organizations comprising at least first and second organizations, and wherein the hub processor is adapted to classify the user according to an organizational affiliation of the user so as to control access by users in the first organization to the data sources held by the second organization.

70. The apparatus according to any of claims 67-69, wherein the hub processor is adapted to receive a query from a user to access the data in the data sources, to determine a query plan for responding to the query by selecting one or more of the data sources responsively to the ontology such that the access rights permit the user to access the data in the one or more of the data sources, and to generate a response to the query in accordance with the query plan.

71. Apparatus for data access, comprising a hub processor, which is adapted to receive a definition of a set of data resources providing access to data, and to collect information regarding a topology and performance characteristics of the data resources, and which is further adapted, upon receiving a query from a user regarding the data, to determine a query plan responsively to the query and to the information regarding the topology and performance characteristics, and to generate a response to the query in accordance with the query plan.

72. The apparatus according to claim 71, wherein the hub processor is adapted to receive measurements of respective load levels of two or more of the data resources, and to select one of the data resources so as to balance the load levels.

73. The apparatus according to claim 71 or 72, wherein the data resources are distributed among a set of autonomous organizations comprising at least first and second organizations, wherein the user submitting the query belongs to the first organization, and wherein the hub processor is adapted to select, responsively to the performance characteristics, one of the data resources of the second organization for use in responding to the query.

74. Apparatus for exchange of information, comprising a plurality of computing nodes, which comprise at least first and second nodes, and which are linked to communicate over a virtual private network running over an underlying public physical network, and which are configured to exchange data packets over the virtual private network in accordance with a semantic communication model, which is defined responsively to an ontology of the information, wherein at least the second node is adapted, upon receiving a data packet over the virtual private network from the first node, to filter the data packet against the semantic communication model so as to verify that the data packet is legitimate.

5 75. The apparatus according to claim 74, wherein the semantic communication model defines a closed set of semantic messages that may be carried by data packets in the virtual private network.

10 76. The apparatus according to claim 74, wherein the nodes are distributed among a set of autonomous organizations.

15 77. The apparatus according to any of claims 74-76, wherein the nodes comprise gateway nodes, which are configured to communicate with clients and data sources using native data formats, and to translate the native data formats to the semantic communication model for communication over the virtual private network.

20 78. The apparatus according to claim 77, wherein the nodes further comprises hub nodes, wherein the gateway nodes are configured to communicate with the hub nodes using the semantic communication model.

25 79. A computer software product, comprising a computer-readable medium in which program instructions are stored, which instructions, when read by a computer, cause the computer to receive a definition of an ontology for application to a set of diverse data sources comprising data having predefined semantics, and to associate with the ontology one or more logical rules applicable to the semantics of the data in the data sources, and further cause the computer, upon receiving a query from a user regarding the data, to determine a query plan for responding to the query by selecting one or more of the data sources responsively to the ontology and by identifying an operation to be applied to the data responsively to the applicable logical rules, and to generate a response to the query in accordance with the query 30 plan.

80. The product according to claim 79, wherein the logical rules comprise a validation rule, and wherein the query plan comprises validating the data from at least one of the data sources responsively to the validation rule.

81. The product according to claim 79, wherein the logical rules comprise a mapping rule, such that at least one of the data sources is mapped to the ontology in accordance with the mapping rule, and wherein the query plan comprises determining an applicability of the at least one of the data sources to the query responsively to the mapping rule.

82. The product according to claim 79, wherein the logical rules comprise a joining rule, and wherein the query plan comprises selecting a key responsively to the joining rule, and joining the data from two or more of the data sources using the key.

83. The product according to claim 82, wherein selecting the key comprises analyzing the data so as to select one or more fields in the two or more of the data sources for use as the key so as to provide a desired statistical probability that the data will be joined correctly.

84. The product according to claim 79, wherein the logical rules comprise a transformation rule, and wherein the query plan comprises transforming the data in at least one of the data sources from a first value that is held in the at least one of the data sources to a second value responsively to the transformation rule.

85. The product according to claim 79, wherein the logical rules comprise a business logic rule, and wherein the query plan comprises processing the data from at least one of the sources responsively to the business logic rule.

86. The product according to claim 79, wherein the logical rules comprise an access rule, and wherein the query plan comprises selecting at least one of the data sources for use in generating the response responsively to the access rule as applied to the user who submitted the query.

87. The product according to any of claims 79-86, wherein the instructions further cause a respective wrapper to be associated with each of the data sources, so as to transform the data from each of the data sources from a native format to an ontological format determined by the ontology, and cause the respective wrapper to apply the operation to the data, and to report the data to the hub processor following application of the operation.

88. The product according to claim 87, wherein the operation applied by the wrapper comprises joining the data from two or more of the data sources.
89. The product according to claim 87, wherein the operation applied by the wrapper comprises mapping values of the data.
- 5 90. The product according to claim 89, wherein mapping the values comprises normalizing the data from a native representation to an ontological representation.
91. The product according to claim 87, wherein the query plan comprises a group of sub-queries, and wherein the instructions cause the computer to send the sub-queries to respective wrappers of a plurality of the data sources, and to combine the data reported from the wrappers
- 10 in order to produce the response.
92. The product according to claim 91, wherein the instructions cause the computer to invoke two or more of the wrappers to operate in parallel.
93. The product according to claim 87, wherein the instructions cause the wrapper to distribute an advertisement of each of the data sources in accordance with the ontology, and
- 15 wherein the instructions cause the computer to discover each of the data sources responsively to the advertisement, and to build the query plan based on the discovered data sources.
94. The product according to claim 87, wherein the instructions cause the wrapper to report the data by sending data packets over a network, the packets comprising semantic content in a form determined by the ontology, and wherein the instructions cause the computer to verify
- 20 legitimacy of the packets responsively to the semantic content.
95. The product according to claim 87, wherein the instructions cause the wrapper to stream the data from the wrapper to a specified storage location.
96. The product according to claim 87, wherein the instructions cause the wrapper to move the data in a block operation from the wrapper to a specified storage location.
- 25 97. The product according to any of claims 79-86, wherein the instructions cause the computer to collect information regarding a topology and performance characteristics of the data sources, and to select, responsively to the information, the data sources to be used in responding to the query.

98. A computer software product, for providing a user with access to a set of diverse information resources, which are configured to provide information with predefined semantics, the product comprising a computer-readable medium in which program instructions are stored, which instructions, when read by a computer, cause the computer to receive a definition of an ontology for application to the information provided by the set of diverse information resources and to associate with the ontology one or more logical rules applicable to the semantics of the information, and further cause the computer, upon receiving a request from the user to access the information, to determine a plan for responding to the request by selecting one or more of the information resources responsively to the ontology and by identifying an operation to be applied to the information responsively to the applicable logical rules, and to generate a response to the request in accordance with the plan.

5 99. The product according to claim 98, wherein the request comprises a query for data held by the information resources.

100. The product according to claim 98, wherein the request comprises a subscription 15 request, and wherein the instructions cause the computer to provide a succession of responses over a period of time responsively to the subscription request.

101. The product according to claim 98, wherein the information resources are configured to provide Web services, and wherein the request specifies one or more of the Web services to be provided to the user.

20 102. The product according to claim 98, wherein the request specifies data to be written to a data repository associated with one or more of the information resources.

103. The product according to claim 98, wherein the request specifies a transaction to be carried out and recorded by one or more of the information resources.

25 104. The product according to any of claims 98-103, wherein the request specifies a process to be carried out by one or more of the information resources.

105. The product according to claim 104, wherein the request specifies an event, and wherein the hub is adapted to cause the specified process to be carried out responsively to an occurrence of the event.

106. A computer software product, comprising a computer-readable medium in which program instructions are stored, which instructions, when read by a computer, cause the computer to receive a definition of an ontology for application to a set of diverse data sources comprising data and a definition of data access rights with respect to the ontology, and to control user access to the data responsively to the ontology of the data and the access rights applicable thereto.

5 107. The product according to claim 106, wherein the ontology comprises a user ontology, and wherein the instructions cause the computer to define the data access rights by assigning a classification to a user according to the user ontology, and to compare the classification to the access rights applicable to the data in order to control the user access.

10 108. The product according to claim 107, wherein the diverse data sources are distributed among a set of autonomous organizations comprising at least first and second organizations, and wherein the instructions cause the computer to classify the user according to an organizational affiliation of the user so as to control access by users in the first organization to 15 the data sources held by the second organization.

20 109. The product according to any of claims 106-108, wherein the instructions cause the computer to receive a query from a user to access the data in the data sources, to determine a query plan for responding to the query by selecting one or more of the data sources responsively to the ontology such that the access rights permit the user to access the data in the one or more of the data sources, and to generate a response to the query in accordance with the 25 query plan.

110. A computer software product, comprising a computer-readable medium in which program instructions are stored, which instructions, when read by a computer, cause the computer to receive a definition of a set of data resources providing access to data, and to collect information regarding a topology and performance characteristics of the data resources, and further cause the computer, upon receiving a query from a user regarding the data, to determine a query plan responsively to the query and to the information regarding the topology and performance characteristics, and to generate a response to the query in accordance with the query plan.

111. The product according to claim 110, wherein the instructions cause the computer to receive measurements of respective load levels of two or more of the data resources, and to select one of the data resources so as to balance the load levels.

112. The product according to claim 110 or 111, wherein the data resources are distributed among a set of autonomous organizations comprising at least first and second organizations, wherein the user submitting the query belongs to the first organization, and wherein the instructions cause the computer to select, responsively to the performance characteristics, one of the data resources of the second organization for use in responding to the query.

113. A computer software product, comprising a computer-readable medium in which program instructions are stored, which instructions, when read by a group of computing nodes that includes at least first and second nodes, linked to communicate over a physical public network, cause the computing nodes to communicate in a virtual private network by exchanging data packets over the public physical network in accordance with a semantic communication model, which is defined responsively to an ontology of the information, wherein the instructions cause at least the second node, upon receiving a data packet over the virtual private network from the first node, to filter the data packet against the semantic communication model so as to verify that the data packet is legitimate.

114. The product according to claim 113, wherein the semantic communication model defines a closed set of semantic messages that may be carried by data packets in the virtual private network.

115. The product according to claim 113, wherein the nodes are distributed among a set of autonomous organizations.

116. The product according to any of claims 113-115, wherein the nodes comprise gateway nodes, wherein the instructions cause the gateway nodes to communicate with clients and data sources using native data formats, and to translate the native data formats to the semantic communication model for communication over the virtual private network.

117. The product according to claim 116, wherein the nodes further comprises hub nodes, wherein the instructions cause the gateway nodes to communicate with the hub nodes using the semantic communication model.